Quasi-Probabilistic Readout Correction of Mid-Circuit Measurements for Adaptive Feedback via Measurement Randomized Compiling

Akel Hashim,[1,][*] Arnaud Carignan-Dugas,[2,][*] Larry Chen,[1] Christian Jünger,[1,3]
Neelay Fruitwala,[3] Yilun Xu,[3] Gang Huang,[3] Joel Wallman,[2] and Irfan Siddiqi[1,3,4]

[1]*Quantum Nanoelectronics Laboratory, Department of Physics,*
University of California at Berkeley, Berkeley, CA 94720, USA
[2]*Keysight Technologies Canada, Kanata, ON K2K 2W5, Canada*
[3]*Computational Research Division, Lawrence Berkeley National Lab, Berkeley, CA 94720, USA*
[4]*Materials Sciences Division, Lawrence Berkeley National Lab, Berkeley, CA 94720, USA*
(Dated: December 22, 2023)

Quantum measurements are a fundamental component of quantum computing. However, on modern-day quantum computers, measurements can be more error prone than quantum gates, and are susceptible to non-unital errors as well as non-local correlations due to measurement crosstalk. While readout errors can be mitigated in post-processing, it is inefficient in the number of qubits due to a combinatorially-large number of possible states that need to be characterized. In this work, we show that measurement errors can be tailored into a simple stochastic error model using randomized compiling, enabling the efficient mitigation of readout errors via quasi-probability distributions reconstructed from the measurement of a single preparation state in an exponentially large confusion matrix. We demonstrate the scalability and power of this approach by correcting readout errors without the need for any matrix inversion on a large number of different preparation states applied to a register of a eight superconducting transmon qubits. Moreover, we show that this method can be extended to measurement in the single-shot limit using quasi-probabilistic error cancellation, and demonstrate the correction of mid-circuit measurement errors on an ancilla qubit used to detect and actively correct bit-flip errors on an entangled memory qubit. Our approach paves the way for performing an assumption-free correction of readout errors on large numbers of qubits, and offers a strategy for correcting readout errors in adaptive circuits in which the results of mid-circuit measurements are used to perform conditional operations on non-local qubits in real time.

Keywords: Quantum Computing, Quantum Measurement, Randomized Compiling

I. INTRODUCTION

Measurement plays a foundational role in quantum mechanics. It is the means by which we learn properties of quantum systems, and is fundamentally linked with the collapse of quantum wavefunctions. Measurement is also essential to quantum computing. In gate-based quantum computing, measurement is needed to translate quantum bit (qubits) to classical bits at the end of a computation, is the central component in teleportation-based protocols [1, 2] and measurement-based quantum computing [3, 4], can be utilized to generate long-range entanglement in constant depth via adaptive quantum circuits [5], and is necessary for syndrome extraction in quantum error correction [6–11]. However, measurements are inherently noisy, and the nature of errors can depend not only on the quantum state prior to measurement, but can also contextually depend on the state of other qubits. Moreover, measurements are often slower and more error prone than the unitary gates used to prepare quantum states, which places limits on the speed and fidelity with which they can be used to perform real-time corrections in the middle of quantum circuits.

A common assumption in quantum computing is that readout errors are purely probabilistic — for a given projective measurement of some finite duration, a qubit has a defined probability of experiencing a bit flip during readout. However, this assumption is often violated in systems with multiplexed readout in which measurement crosstalk can cause context-dependent and correlated readout errors [12]. Moreover, the probability of a bit-flip for a qubit in an excited state can vastly differ from the probability of a bit-flip while sitting in the ground state due to non-unital processes such as T_1 decay. However, by twirling a process over a unitary 1-design, one can effectively *design* stochastic channels [13]. One such strategy for designing stochastic channels is randomized compiling (RC) [14, 15], which is a robust and efficient method for tailoring arbitrary Markovian errors into Pauli channels in gate-based quantum computing. While RC was originally designed for tailoring gate noise, it can be adapted to tailor measurement noise [16], and has been previously shown to reduce worst-case error rates in state-preparation and measurement (SPAM) [17]. In this work, we experimentally deploy RC for quantum measurements on an eight-qubit superconducting quantum processor (see Fig. 1a). We show that, under measurement RC (MRC), quantum measurement noise can be accurately described by a stochastic error model in which the probability of a bit flip for any given qubit is independent of the preparation state or

the state of other qubits on the quantum processor, thus enforcing common pre-existing assumptions about the stochasticity of measurement errors.

Because measurements translate quantum bits to classical bits, errors in measurements can be corrected classically via post-processing. To do so generally requires preparing and measuring all possible combinations of input basis states for n qubits, from which one can construct a *confusion matrix* of the measured results. This confusion can be inverted to correct to readout errors, but the size of this matrix grows exponentially in the number of qubits, making both the characterization and inversion steps intractable for large qubit numbers. As a result, experimentalists often resort to perform *local* readout correction [18], in which an individual confusion matrix is measured and inverted for each qubit. While this can correct individual readout errors, it cannot correct correlated bit-flips. By tailoring noise in measurement into a stochastic bit-flip channel, we show that it is possible to correct readout errors for any input state without any matrix inversion. To do so, it is sufficient to characterize a single input states (e.g., $|0^{\otimes n}\rangle$ for n qubits) under MRC, from which a quasi-probability distribution can be constructed. Readout correction is then performed by inverting the quasi-probability distribution on the measured bit-string results. We compare local readout correction to our quasi-probabilistic protocol for a large number of structured and random input states on eight qubits, and show that our protocol improves the results in 98.5% of the circuits. Moreover, we show that this scheme extends to mid-circuit measurements (MCMs), and demonstrate the mitigation of readout errors used to perform real-time feedback to correct for bit-flip errors on an entangled qubit, improving its effective lifetime.

II. RANDOMIZED COMPILING FOR MEASUREMENTS

Generalized measurements of quantum states are described by positive-operator valued measures (POVMs), which are set of positive semi-definite Hermitian matrices $\{E_i\}$ in d-dimensional Hilbert space $\mathcal{H}_d$ that obey the completeness relation:

$$\sum_i E_i = \mathbb{I}. \tag{1}$$

The probability of measuring outcome E_i given a state ρ is governed by Born's rule,

$$p(i|\rho) = \text{Tr}[E_i \rho]. \tag{2}$$

For a given a system containing n qubits, the POVM corresponding to computational basis measurements contains 2^n elements, $\{E_i\}_{i=1}^{2^n}$, with each element indexed by an n-qubit bit string i. For example, for a single qubit the POVM set is $\{E_0, E_1\}$, for two qubits the POVM set is $\{E_{00}, E_{01}, E_{10}, E_{11}\}$, etc. By preparing a system of n qubits in all 2^n possible combinations of basis states, represented by the set of input states $\{\rho_j\}$, and measuring the resulting POVMs $\{E_i\}$ for each basis state, one can construct a $2^n \times 2^n$ *confusion matrix* $\mathcal{M} = \langle\langle\{E_i\}|\{\rho_j\}\rangle\rangle$ whose elements

$$\mathcal{M}_{ij} = \text{Tr}[E_i \rho_j] \tag{3}$$

represent the probability $p(i|j)$ of measuring the outcome E_i given an input state ρ_j. In general, quantum measurement errors cannot be modeled by a confusion matrix because it ignores the effect of the measurement process on quantum superpositions.

However, let's assume for a moment that the nature of the measurement error is such that given and n-qubit confusion matrix $\mathcal{M}$ and an ideal probability distribution p, the effect of measurement noise on the ideal outcomes produces in a noisy probability distribution $q = \mathcal{M}p$. In such case, correcting the effect of measurement noise on a probability distribution reduces to inverting $\mathcal{M}$ given a measured distribution q:

$$p = \mathcal{M}^{-1}q. \tag{4}$$

If $\mathcal{M}$ is known and if it correctly models measurement errors, then in theory one can correct the effect of measurement errors affecting the outcome of any quantum algorithm. However, because $\mathcal{M}$ scales exponentially in the number of qubits n, in practice it is not feasible a full n-qubit confusion matrix, nor is it always necessary if one can make reasonable assumptions about the locality and nature of correlated measurement noise. An alternative strategy is assume readout errors are uncorrelated and that measurement noise can be modeled as a tensor product of confusion matrices. In this case, it is sufficient to reconstruct the individual confusion matrix for each qubit, such that $\mathcal{M}$ is given as

$$\mathcal{M} = \prod_{i=0}^{n} \otimes \mathcal{M}_i, \tag{5}$$

where $\mathcal{M}_i$ is the confusion matrix for the ith qubit. Now, the inversion process (Eq. 4) only corrects readout errors on each qubit individually, but cannot account for any correlated readout errors. While it is often assumed that readout errors are probabilistic and locally independent, in which case measuring individual confusion matrices for each qubit would be sufficient to correct all readout errors, in practice this is not the case. For example, in Fig. 1a, we plot the full confusion matrix for three qubits (Q0, Q1, and Q2; see Fig. 1a and Appendix). We observe that for most preparation states, the readout fidelity is between $85\% - 95\%$. However, for $|010\rangle$ and $|110\rangle$ we observe poor readout fidelities and large state-dependent errors. Such errors are likely due in part to the fact that the readout frequencies for Q0 and Q1 are close in frequency resulting from fabrication defects (see Appendix A), leading to readout crosstalk. Moreover, even for the preparation states with higher readout fidelities, we
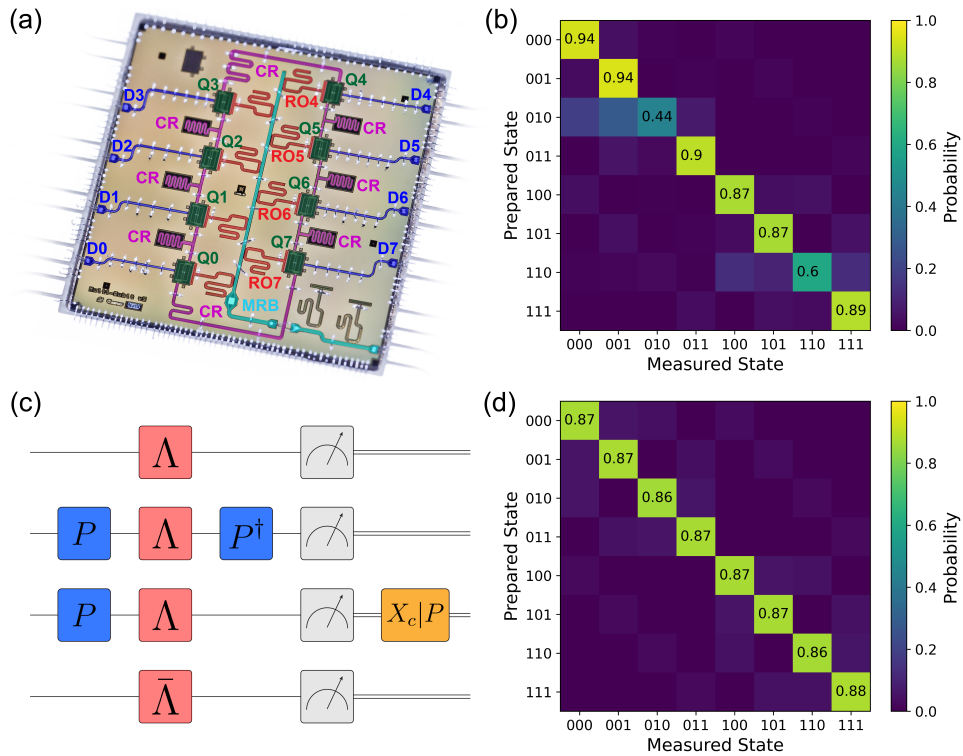
Figure 1: Randomized Compiling for Measurements. (a) 8-qubit superconducting transmon processor. Qubits are labeled in green, individual drive lines are labeled in blue, individual readout resonators (RO) are labeled in red, and the multiplexed readout bus (MRB) is labeled in cyan. The qubits are coupled to nearest neighbors in a ring geometry via coupling resonators (CR, purple). (b) Full confusion matrix measured for three qubits (Q0, Q1, Q2). Strong state-dependent errors are observed. For example, when the $|010\rangle$ state is prepared, $|000\rangle$ is measured 19% of the time, and $|001\rangle$ is measured 28% of the time. (c) [Top] We can model the error in a measurement by a process matrix Λ preceding the measurement. [Second] In theory, it is possible to twirl this process matrix via Pauli twirling, $\Lambda \mapsto P\Lambda P^\dagger$. [Third] However, because this is a process matrix for a (non-unitary) measurement, the inversion operators must be implement as classical bit flips X_c conditioned on which Pauli P was sampled before the measurement. [Bottom] By averaging this measurement many times over the full Pauli group, we obtain a twirled error process $\bar{\Lambda} = \sum_{P \in \{I,X,Y,Z\}} P\Lambda P^\dagger$, in which measurement errors have been reduced to a stochastic bit-flip channel. (d) Full three-qubit confusion matrix measured using the scheme presented in (c) for qubits Q0, Q1, Q2. We observe that the diagonal entries of the confusion matrix are all approximately equal in magnitude, and all error probabilities in the off-diagonal elements are symmetric. This indicates that, under randomized compiling, the probability of a bit flip for any qubit is the same, regardless of the state that is prepared.

generally observe that excited states have worse readout fidelity than ground states. This is due to non-unital errors such as T_1 decay, which places fundamental limits on excited state fidelities for a given readout time.

One way to ensure that measurement errors can be modeled by a confusion matrix and do not exhibit state- and context-dependent features is to twirl the readout noise into stochastic bit-flip channels using methods such as randomized compiling for quantum measurements [16]. To do so, we model a noisy measurement $\langle\langle \tilde{E}_i|$ as an ideal measurement $\langle\langle E_i|$ preceded by a process matrix Λ which captures all measurement errors: $\langle\langle \tilde{E}_i| = \langle\langle E_i| \Lambda$ (see Fig. 1b). The goal of measurement randomized compiling (MRC) is to twirl Λ into diagonal Pauli channels, i.e., $\Lambda \mapsto \sum_{P \in \mathbb{P}_n} P^\dagger \Lambda P$, where $\mathbb{P}_n = \{I, X, Y, Z\}^{\otimes n}$ is the n-qubit Pauli group. However, in reality readout errors occur concurrently with measurement;

therefore, we cannot simply conjugate Λ by Pauli gates. Rather, to twirl Λ we compile random Paulis into the final cycle of single-qubit gates before measurement and perform classical bit-flips on the measured results conditional on the inserted Pauli for each qubit. For example, if I or Z is inserted, these will not change the results of measurements in the computational basis; however, if X or Y is chosen, these will flip the qubit state prior to measurement, necessitating classical bit-flips after measurement. By repeating this process many (K) times and recording the combined distribution of all results, we obtain an effective Pauli-twirled process matrix

$$\bar{\Lambda} = \sum_{P \in_R \mathbb{P}_n}^{K} P^\dagger \Lambda P, \qquad (6)$$

where R denotes that P is chosen at random from the n-qubit

(a)

$$\mathcal{F} = 75.5\%$$

(b)

$$\mathcal{F} = 82.4\%$$
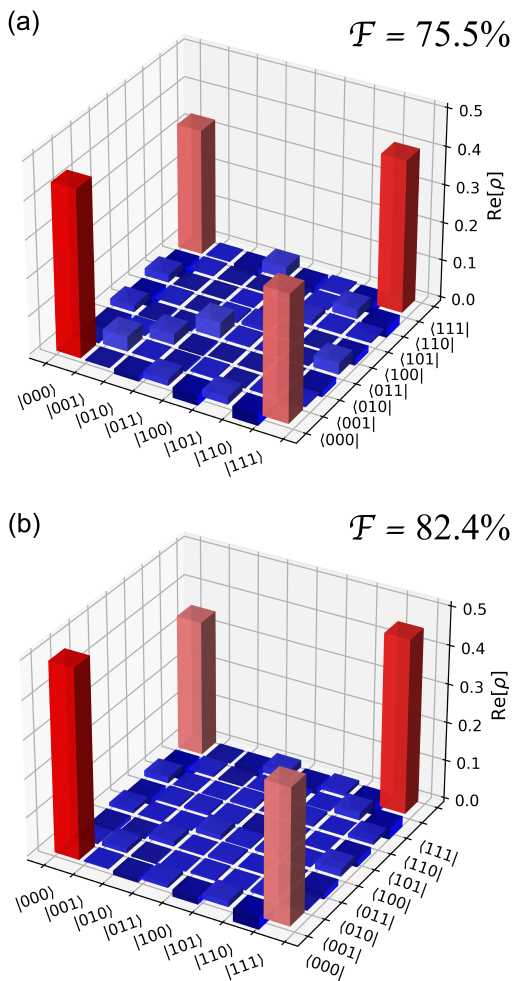
Figure 2: Correcting Readout Errors in Expectation Values. (a) Density matrix of a three-qubit GHZ state reconstructed via state tomography. The raw measured state fidelity is 75.5%. (b) Density matrix of a three-qubit GHZ state measured with MRC and reconstructed via state tomography after re-scaling each expectation value by the average readout fidelity given by the confusion matrix in 1c. The measured state fidelity is 82.4%, with the remaining error likely due to decoherence and gate errors.

Pauli group $\mathbb{P}_n$ each time. In Fig. 1c, we plot the full confusion matrix for qubits Q0, Q1, and Q2 reconstructed using MRC with $K = 100$ randomizations [17]. We observe that the diagonal readout fidelities $p(i|i)$ are all approximately equal, showing that we have eliminated state-dependent errors due to T_1 decay. Moreover, the off-diagonal probabilities $p(i|j) \ \forall \ j \neq i$ are symmetric along the diagonal, suggesting that we have eliminated context-dependent readout errors due to readout crosstalk. Therefore, under MRC, we can describe readout errors as a purely stochastic process in which the probability of a bit flip for any given qubit is independent of the preparation state. A similar method was introduced in [19] by inserting random bit-flips prior to

measurement. However, bit-flip averaging does not provide a complete twirl of the readout noise, as phase randomization is also necessary in order to describe readout errors are purely stochastic. For example, suppose a qubit is in the $|i+\rangle$ state prior to measurement; here, a coherent-X error during measurement will result in an incorrect results distribution. However, by randomly insert Pauli-Z gates prior to measurement, the impact of the coherent-X error will be averaged away on the ensemble level.

Because readout errors are state-independent under MRC, the effect of readout errors on expectation values can be efficiently corrected by re-scaling by the average readout fidelity. To demonstrate this, we perform state tomography on a three-qubit GHZ state with and without MRC, shown in Fig. 2. When performing state tomography without MRC (Fig. 2a), we measure a state fidelity of 75.5% (Fig. 2a). When performing state tomography with MRC (fig2a), we measure a raw state fidelity of 73.2%; however, when we re-scale each expectation value by the average measurement fidelity given by the confusion matrix in Fig. 1c, the fidelity improves to 82.4%. The remaining infidelity is likely due to decoherence and gate errors. While this re-scaling is efficient for expectation values, to correct readout errors at the level of raw bit-strings, we must invert the probability of bit-flip errors for each qubit. In the following section, we introduce a method for doing so based on quasi-probabilities.

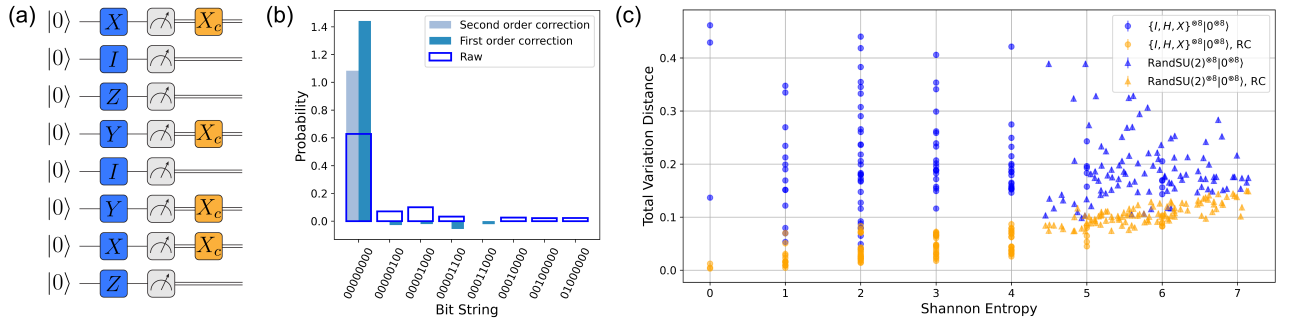## III. QUASI-PROBABILISTIC READOUT CORRECTION

As observed in the previous section, applying randomized compiling to quantum measurements effectively tailors the measurement error channel Λ into a classical stochastic error channel $\bar{\Lambda}$:

$$\bar{\Lambda}[\rho] = \sum_{x \in \mathbb{Z}_m^{\otimes n}} p_x X_x \rho X_x , \qquad (7)$$

where $x \in \mathbb{Z}_2^{\otimes n}$ is the set of classical n-bit strings, $\{p_x\}_{x \in \mathbb{Z}_2^n}$ is a probability distribution over bit-flips X_x, and where X_x is short for $X_{x_1} X_{x_2} \cdots X_{x_n}$. This has a few ramifications: firstly, the effective error channel $\bar{\Lambda}$ can be fully described by its corresponding probability distribution, and each probability p_i can be estimated up to $1/\sqrt{N_{\text{shots}}}$ simply by looking at the output distribution resulting from sending a single computational basis input to the randomly compiled measurement channel. In other words, $\bar{\Lambda}$ can be approximately described with $O(N_{\text{shots}})$ floating-point numbers, and each number has a precision of $1/\sqrt{N_{\text{shots}}}$. [1] Secondly, the effective mea-

---

[1] The accuracy of the estimate is also limited by state preparation errors and single-qubit gate errors, but these tend to be low compared to measurement errors.

Figure 3: Quasi-Probabilistic Readout Correction. (a) Readout characterization. The probability of bit-flip errors during readout can be characterized by preparing a single n-qubit input state (e.g. $|000000000\rangle$ for eight qubits) and measuring the resultant states under RC. To do so, a randomly-sampled n-qubit Pauli operator should be inserted before measurement (e.g. $X \otimes I \otimes Z \otimes Y \otimes I \otimes Y \otimes X \otimes Z$); after measurement, a classical bit-flip X_c should be applied to any qubit in which an X or Y gate was applied before measurement. This process should be repeated many times for many different randomly-sampled Pauli operators. (b) Results for the protocol in (a) applied to eight qubits using $N = 100$ different randomizations is plotted in blue. We observe significant single-qubit bit-flips errors on many qubits. For example, while $|000000000\rangle$ is measured over 60% of the time, we observe that qubit 4 has a 10% chance of experiencing a bit-flip error during readout. This distribution can be used to perform quasi-probabilistic readout correction on any 8-qubit circuit in which the readout is performed using measurement RC. The first and second order corrections performed on the raw distribution are plotted in purple and orange, respectively. The second order correction reconstructs a distribution in which only the all zero state remains. (Only the bit strings with significant counts are included for clarity.) (c) Readout-corrected single-qubit circuits with and without measurement RC. Structured were generated by randomly sampling from $\{I, H, X\}$ gates for each qubits (circular data points), and random circuits were generated by applying random SU(2) gates independently to each qubit (triangular data points). Each circuit was performed with and without MRC, denoted by the orange and blue data points, respectively. For circuits without MRC, we apply local readout correction using confusion matrices measured for each qubit. We plot the TVD of the experimental results with the ideal results as a function of the Shannon entropy of the ideal result; the larger the entropy, the more uniform the distribution. We observe that the MRC results broadly outperform the results with local readout correction, with better performance at lower entropy. (Error bars on the TVD are on the order of the size of the markers.)

surement error $\bar{\Lambda}$ can be inverted by applying a linear operation on the noisy output distribution. The exact inversion can quickly become unscalable to describe, but since the probabilities appearing in $\bar{\Lambda}$ are already estimated with $1/\sqrt{N_{\text{shots}}}$ precision, an approximation should suffice. Fortunately, there exist standard quasi-probabilistic correction techniques that provide different orders of approximation of the inverse of $\bar{\Lambda}$. The first order approximation is described using $O(N_{\text{shots}})$ floating-points numbers, and in general the ith order approximation is described using $O(N_{\text{shots}}^i)$ floating-points numbers. We proceed with describing the quasi-probabilistic readout correction (QPRC) protocol below.

Because measurement errors under MRC can be described by a stochastic bit flip channel which is independent of the input state, it is sufficient to characterize the probability of bit-flips on n qubits using a single preparation state. For simplicity, we choose to characterize measurement errors on $|0^{\otimes n}\rangle$ using the MRC protocol. For example, in Fig. 3a we depict a single cycle of Paulis applied to the all-zero state on eight qubits; if X or Y is applied before measurement, then a classical bit-flip X_c is applied in post-processing. This process should be repeated many (K) times to construct a twirled measurement channel (Eq. 7). In Fig. 3b, we plot the results of the characterization procedure in blue using $K = 100$ randomizations. We observe that the all-zero

state is measured over 60% of the time, with the remaining 40% distributed over various single- and multi-qubit bit-flip channels. Now, given a characterized error probability distribution $p = \sum_x p_x x$ and an ideal outcome distribution $I = \sum_y I_y y$, we can express the resulting noisy outcome distribution $N = \sum_z N_z z$ as

$$N = \left(\sum_x p_x x\right) \oplus \left(\sum_y I_y y\right) \tag{8}$$

where $y \oplus x$ is the *bitwise* modulo 2 sum of the y and x bit-strings. In other words, the probability of observing the outcome z given a noisy measurement is

$$N_z = \sum_{x \oplus y = z} p_x I_y . \tag{9}$$

There are many possible strategies to invert the effect of readout errors on the distribution I. The simplest one is to construct a quasi-probability distribution q which is an approximate inverse of p (with regards to $\oplus$):

$$q = \frac{1}{2p_{\mathbf{0}} - 1}\left(p_{\mathbf{0}}\mathbf{0} - \sum_{x \neq \mathbf{0}} p_x x\right), \tag{10}$$

where $\mathbf{0}$ is shorthand for $0^{\otimes n}$. Indeed, applying q yields

$$p' = q \oplus p = \frac{1}{2p_\mathbf{0} - 1}\left(p_\mathbf{0}\mathbf{0} - \sum_{x \neq \mathbf{0}} p_x x\right) \oplus \left(p_\mathbf{0}\mathbf{0} + \sum_{x \neq \mathbf{0}} p_x x\right)$$

$$= \frac{p_\mathbf{0}^2}{2p_\mathbf{0} - 1}\mathbf{0} - \frac{1}{2p_\mathbf{0} - 1}\left(\sum_{x \neq \mathbf{0}} p_x x\right)^2 . \qquad (11)$$

To put it simply, the error amplitude goes from $1 - p_\mathbf{0}$ to $(1 - p_\mathbf{0})^2/(2p_\mathbf{0} - 1)$. The inverse operation can be improved. Indeed, consider the family of quasi-probability distributions:

$$q^{(k)} = \frac{p_\mathbf{0}^{2k-1}}{p_\mathbf{0}^{2k} - (1 - p_\mathbf{0})^{2k}}\left(\mathbf{0} + \sum_{j=1}^{2k-1}\left(\frac{-1}{p_\mathbf{0}}\right)^j\left(\sum_{x \neq \mathbf{0}} p_x x\right)^j\right), \quad (12)$$

for $k \in \mathbb{N}_+$. Notice that $q^{(1)} = q$ from Eq. (10). Applying $q^{(k)}$ to p yields:

$$q^{(k)} \oplus p = \frac{p_\mathbf{0}^{2k}}{p_\mathbf{0}^{2k} - (1 - p_\mathbf{0})^{2k}}\mathbf{0} - \frac{1}{p_\mathbf{0}^{2k} - (1 - p_\mathbf{0})^{2k}}\left(\sum_{x \neq \mathbf{0}} p_x x\right)^{2k} . \qquad (13)$$

In this generalized case, the error amplitude goes from $1 - p_\mathbf{0}$ to $\frac{1}{1 - (p_\mathbf{0}/(1-p_\mathbf{0}))^{2k}}$.

In practice, to correct readout errors on any noisy experimental probability distribution N which has been measured using MRC, we sum over all corrected results in which the counts for each experimental results N_x have been redistributed according to $q^{(k)}$:

$$N^{(k)} = \sum_x \left(N_x x \oplus q^{(k)}\right), \qquad (14)$$

where the readout corrected distribution $N^{(k)}$ is the union over all of the redistributed counts $N_x x \oplus q^{(k)}$.

To demonstrate that our procedure corrects readout errors, we perform a first (i.e. using $q^{(1)} = q$) and second (i.e. using $q^{(2)}$) order correction on the characterized probability distribution in Fig. 3b that is used to construct the quasi-probability distribution. We observe that the first correction redistributes most of the results to $\mathbf{0}$, but that there remain significant (negative) quasi-probabilities in other bins. Because the corrected distribution is itself a quasi-probability distribution that has been normalized to preserve the total probability, $\mathbf{0}$ has a quasi-probability $p_\mathbf{0}^{(1)}$ greater than one to account for the negative quasi-probabilities in the other states. After performing a second order correction on the characterized distribution, we find that $p_\mathbf{0}^{(2)} \approx 1$, as we would expect if all readout errors were corrected. This process highlights the fact that readout correction (or, more generally, error mitigation strategies) can introduce non-physical outcomes into the results of experiments. For example, if one wants to preserve the total probability of a process, then the small residual negative values that remain after

the quasi-probabilistic error correction should be preserved, which equates to enforcing trace-preservation (TP). However, negative probabilities violate complete-positivity (CP), and these values could be reasonably set to zero depending on the nature of the final computation. Therefore, in general one cannot enforce both CP and TP on the outcomes of error corrected results, and the choice of which to preserve is up to the experimenter.

To demonstrate the efficacy of our protocol on a wide variety of input states, we perform a second-order correction (i.e. using $q^{(2)}$) on 200 different eight-qubit circuits consisting of a single-cycle of gates, shown in Fig. 3c. For half of the circuits, we sample gates from $\{I, H, X\}$ at random for each qubit, and for the other half of the circuits we sample random SU(2) gates for each qubit independently. To compute the accuracy of the readout corrected results, we compute the total variation distance (TVD) between the experimental distribution N and the ideal distribution I,

$$D_{\mathrm{TV}}(N, I) = \frac{1}{2}\sum_x |N_x - I_x|, \qquad (15)$$

plotted as a function of the Shannon entropy of the ideal results,

$$S = -\sum_x I_x \log_2(I_x). \qquad (16)$$

We compare the results of our QPRC protocol to results obtained using local readout correction, and find that our protocol produces better results in 98.5% of the circuits. Moreover, we observed a positive linear correlation between the TVD of the corrected results and the entropy of the ideal results, with better performance at lower entropy. This could be explained by the fact that for higher entropy, the approximate correction has to be applied to more outputs, meaning that the systematic error in the approximated inverse is applied more often. No correlation in performance is observed for the results obtained with local readout correction, highlighting that the efficacy of local readout correction cannot be trusted for all input states.

IV. READOUT CORRECTION FOR MID-CIRCUIT MEASUREMENTS

The QPRC protocol presented in the previous section provides a clear strategy for correcting readout errors afflicting the final measurements in a quantum circuit. However, it is less clear how to correct readout errors in mid-circuit measurements (MCM), whose results can be used to adapt circuits in real-time via classical feedback [5]. While the results of a single measurement used for decision branching in feedforward schemes cannot be corrected in real-time, the results of MCMs can still be corrected quasi-probabilistically in the paradigm where we still end up with a distribution
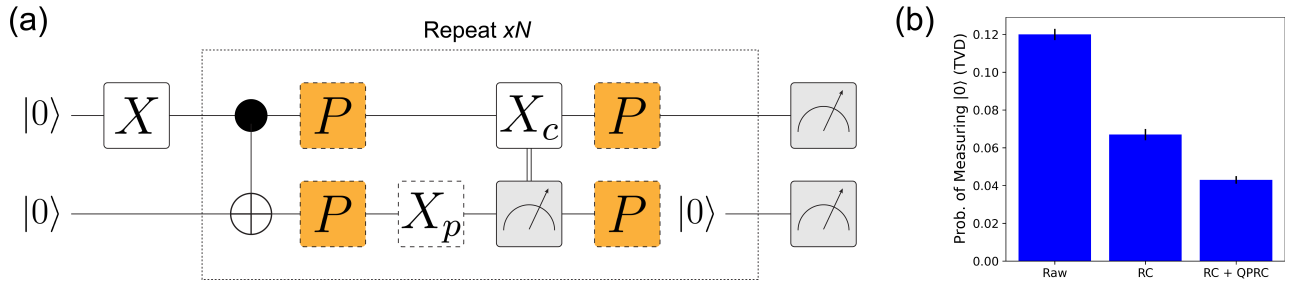
Figure 4: Quasi-Probabilistic Readout Correction of Mid-Circuit Measurements. (a) Schematic for active bit-flip protection. A memory qubit [top] is prepared in the $|1\rangle$ state and entangled with an ancilla qubit [bottom], which is subsequently measured. A conditional bit-flip (X_c) is performed on the memory qubit depending on the results of the measurement on the ancilla qubit, after which the ancilla qubit is reset. This process is performed for N repetitions to protect the memory qubit from decaying to the ground state. Under MRC, the MCM is performed with RC by insertion of random Paulis (P, dashed orange boxes) before and after the MCM. To perform quasi-probabilistic readout correction on the MCM, a Pauli-X gate is probabilistically inserted before the measurement (X_p, dashed white box), and the final results of the MCM with X_p are subtracted from the results without X_p. (b) Results from performing the scheme presented in (a) for a single round of bit-flip protection. In the bare case without MRC, the probability of measuring the memory qubit in the $|0\rangle$ state at the end of the circuit is 12(3)%. When the MCM is performed with RC this improves to 6.8(3)%, suggesting that RC is twirling the noise on the idle qubit that occurs during the MCM. When the results measured with X_p are subtracted from the results measured without X_p, the probability of measuring the memory qubit in $|0\rangle$ reduces to 4.3(2)%.

at the very end of a circuit. To do so requires characterizing the probability of bit-flips for a given MCM, and quasi-probabilistically cancelling this error via random insertion of artificial Pauli-X errors. We describe this procedure below.

When MCMs are used to perform conditional feed-forward operations, the readout fidelity of each MCM will dictate the rate at which the incorrect conditional operation is performed, which will add up linearly as a function of the number of MCMs in the circuit. In a model in which readout errors are purely probabilistic, this rate can be measured *a priori* by characterizing the probability of a bit-flip error on the measured qubit(s). For example, suppose a MCM qubit prepared in the ground state has a probability p_1 of experiencing a bit-flip during measurement, then the probability with which a single instance of the MCM performs the correct conditional operation is $1 - p_1 = p_0$. According to the QPRC protocol presented in the previous section, the results of imperfect measurement can be corrected by assigning a negative weight to the incorrect outcomes and subtracting them from the ideal outcomes. To do so in circuits with MCMs, we probabilistically insert artificial bit-flip X_p prior to the measured qubit with probability $p = p_1$. Now, for a circuit measured N_s times, on average the correct conditional operation will have been applied $(1 - p)N_s$ times, and the incorrect conditional operation will have been applied pN_s times. To mitigate the impact of the noisy MCM, we subtract the raw counts of the circuit measured with X_p from the raw counts of the circuit measured without X_p. For circuits with multiple rounds of MCM, we assign a negative weight to each instance in which X_p appears in the circuit; thus, for odd (even) occurrences, the results are subtracted (added) to the bare results. This process naturally increases the shot noise, since the error mitigated results only have $(1 - p)N_s - pN_s = (1 - 2p)N_s$ shots; one can choose to compensate for this at the cost of a larger overhead by increasing the total number of shots to $N'_s = N_s/(1 - 2p)$.

We demonstrate the correction of readout errors on MCMs by performing the above protocol on a circuit designed to protect the memory of a qubit in the $|1\rangle$ state, shown in Fig. 4a. Real-time active feedback is performed using the open-source control hardware QubiC [20]. When MRC is utilized for MCMs, the conditional readout value of the measured qubit now depends on the Pauli that is sampled before readout, and the conditional operation on the memory qubit must take into account the Paulis before and after the MCM. We find that for a single round of bit-flip protection, the probability of measuring the memory qubit in $|0\rangle$ is 12(3)%, given by computing the TVD of the measured results with the ideal result of being in $|1\rangle$. When we perform the MCM with MRC, this probability is reduced to 6.8(3)%. This improvement is likely due to the impact of Pauli twirling the idle qubit during the measurement, similar to dynamical decoupling. When we perform the QPRC on final results, we find that the probability of measuring the qubit in $|0\rangle$ is 4.3(2)%. The difference between RC and RC + QPRC is consistent with a bit-flip rate of 2.4% measured for the ancilla qubit prior to the experiment. The remaining TVD of $\sim 4\%$ is consistent with the T_1 coherence limit of the memory qubit during the readout on the ancilla qubit.

It should be noted that [21] proposes a related method for mitigating Pauli errors that occur druing MCMs using a quasi-probabilistic error cancellation scheme that utilizes randomized compiling, termed Pauli error cancellation (PEC) [22]. There are advantages and disadvantages to both techniques. A key distinction is that the protocol in [21] utilizes cycle benchmarking [23] to characterize the rates of Pauli errors, which has a much higher characterization overhead than our technique, whose characterization over-

head is constant in the number of qubits n. However, PEC can mitigate global errors that occur across an entire register of qubits, including correlated gate errors, whereas QPRC is only designed to mitigate readout errors. Future work could explore scalability and trade-offs between these related methods as they relate to MCMs and adaptive circuits.

V. DISCUSSION

Improving the fidelity of qubit readout is equally as important as improving gates fidelities. However, in recent years much more focus has been placed on improving gate fidelities, leaving readout errors (or more generally SPAM errors) much larger than contemporary gate errors. To compensate for this, experimentalists typical correct readout errors by inverting a $2^n \otimes 2^n$ confusion matrix or, alternatively, inverting local 2×2 readout confusion matrices for each qubit independently. While the former method can correct n-qubit readout errors that occur on computational basis states, it is not scalable; on the other hand, while the latter method is scalable, it cannot correct correlated readout errors.

In this work, we introduce a quasi-probabilistic method for correcting measurement noise which utilizes randomized compiling for enforcing a stochastic bit-flip model of readout errors. Our method is requires a minimal characterization overhead which is constant in the number of qubits, and is scalable in the limit that probabilistically-small readout errors can be ignored. We demonstrate that our method vastly outperforms local readout correction on a large number of different possible input states for eight qubits. Moreover, we show that it can be extended to measurements in the single-shot limit, such as those used for adaptive circuits or quantum error correction (QEC). We show that we can effectively extend the memory of a qubit in an excited state by quasi-probabilistically cancelling the impact of readout errors on an entangled ancilla qubit.

While significant research and development is required to improve the readout fidelities of contemporary qubits on many hardware platforms, scalable, matrix-inversion-free readout correction methods such as QPRC are useful tools for correcting readout errors in the NISQ era and beyond. Our method is full compatible MCMs, and future work could demonstrate the utility of utilizing QPRC for correcting readout errors in adaptive circuits used for preparing non-local entangled states. Furthermore, the machinery needed for adaptive circuits is the same as what is needed for QEC, so combining QPRC with QEC would be an intriguing avenue for exploration.

[1] D. Gottesman and I. L. Chuang, Nature **402**, 390 (1999).
[2] K. S. Chou, J. Z. Blumoff, C. S. Wang, P. C. Reinhold, C. J. Axline, Y. Y. Gao, L. Frunzio, M. Devoret, L. Jiang, and R. Schoelkopf, Nature **561**, 368 (2018).
[3] R. Raussendorf and H. J. Briegel, Physical review letters **86**, 5188 (2001).
[4] H. J. Briegel, D. E. Browne, W. Dür, R. Raussendorf, and M. Van den Nest, Nature Physics **5**, 19 (2009).
[5] M. Foss-Feig, A. Tikku, T.-C. Lu, K. Mayer, M. Iqbal, T. M. Gatterman, J. A. Gerber, K. Gilmore, D. Gresh, A. Hankin, *et al.*, arXiv preprint arXiv:2302.03029 (2023).
[6] P. W. Shor, Physical review A **52**, R2493 (1995).
[7] E. Knill and R. Laflamme, Physical Review A **55**, 900 (1997).
[8] J. Chiaverini, D. Leibfried, T. Schaetz, M. D. Barrett, R. Blakestad, J. Britton, W. M. Itano, J. D. Jost, E. Knill, C. Langer, *et al.*, Nature **432**, 602 (2004).
[9] N. Ofek, A. Petrenko, R. Heeres, P. Reinhold, Z. Leghtas, B. Vlastakis, Y. Liu, L. Frunzio, S. Girvin, L. Jiang, *et al.*, Nature **536**, 441 (2016).
[10] S. Rosenblum, P. Reinhold, M. Mirrahimi, L. Jiang, L. Frun-

zio, and R. J. Schoelkopf, Science **361**, 266 (2018).
[11] W. P. Livingston, M. S. Blok, E. Flurin, J. Dressel, A. N. Jordan, and I. Siddiqi, Nature communications **13**, 2307 (2022).
[12] J. Heinsoo, C. K. Andersen, A. Remm, S. Krinner, T. Walter, Y. Salathé, S. Gasparinetti, J.-C. Besse, A. Potočnik, A. Wallraff, *et al.*, Physical Review Applied **10**, 034040 (2018).
[13] M. A. Graydon, J. Skanes-Norman, and J. J. Wallman, arXiv preprint arXiv:2201.07156 (2022).
[14] J. J. Wallman and J. Emerson, Phys. Rev. A **94**, 052325 (2016).
[15] A. Hashim, R. K. Naik, A. Morvan, J.-L. Ville, B. Mitchell, J. M. Kreikebaum, M. Davis, E. Smith, C. Iancu, K. P. O'Brien, I. Hincks, J. J. Wallman, J. Emerson, and I. Siddiqi, Phys. Rev. X **11**, 041039 (2021).
[16] S. J. Beale and J. J. Wallman, arXiv preprint arXiv:2304.06599 (2023).
[17] A. Hashim, S. Seritan, T. Proctor, K. Rudinger, N. Goss, R. Naik, J. M. Kreikebaum, D. Santiago, and I. Siddiqi, npj Quantum Inf **9** (2023), 10.1038/s41534-023-00764-y.
[18] S. Bravyi, S. Sheldon, A. Kandala, D. C. Mckay, and J. M.

Gambetta, Physical Review A **103**, 042605 (2021).

[19] A. W. Smith, K. E. Khosla, C. N. Self, and M. Kim, Science advances **7**, eabi8009 (2021).

[20] Y. Xu, G. Huang, N. Fruitwala, A. Rajagopala, R. K. Naik, K. Nowrouzi, D. I. Santiago, and I. Siddiqi, arXiv preprint arXiv:2309.10333 (2023).

[21] R. S. Gupta, E. v. d. Berg, M. Takita, K. Temme, and A. Kandala, arXiv preprint arXiv:2310.07825 (2023).

[22] S. Ferracin, A. Hashim, J.-L. Ville, R. Naik, A. Carignan-Dugas, H. Qassim, A. Morvan, D. I. Santiago, I. Siddiqi, and J. J. Wallman, arXiv preprint arXiv:2201.10672 (2022).

[23] A. Erhard, J. J. Wallman, L. Postler, M. Meth, R. Stricker, E. A. Martinez, P. Schindler, T. Monz, J. Emerson, and R. Blatt, Nature communications **10**, 5347 (2019).

[24] P. Ivashkov, G. Uchehara, L. Jiang, D. S. Wang, and A. Seif, to appear (2023).

[25] F. Mallet, F. R. Ong, A. Palacios-Laloy, F. Nguyen, P. Bertet, D. Vion, and D. Esteve, Nature Physics **5**, 791 (2009).

Appendix A: Qubit & Readout Characterization

The quantum processing unit (QPU) used in this work consists of eight superconducting transmon qubits arranged in a ring geometry (Fig. 1a). The frequency spectrum of the GE and EF transition of each qubit is plotted in Fig. A1a. Some frequency crowding is observed at the lower end of the frequency spectrum. For example, the GE transition of Q2 is close to the EF transitions of Q0, Q4, and Q3. These can lead to microwave line crosstalk between qubits, which can result in coherent leakage on the EF transitions when the GE transition of Q2 is driven. A similar effect can occur between the GE transition of Q5 and the EF transition of Q6, which is spectrally far from the rest of the qubits on the QPU due to fabrication inaccuracies. The qubit coherences, randomized benchmarking fidelities, and cycle benchmarking fidelities are listed in Table A1.

In Fig. A1c, we plot the readout calibration for all eight qubits on this QPU, which supports qutrit state discrimination. In qubit computations, qutrit readout can be used to measure leakage rates. Alternatively, qutrit state discrimination can be used for excited state promotion (ESP) [25] for improving qubit readout fidelities, whereby an $\pi_{1\to2}$ pulse is applied to each qubit before readout, after which all $|2\rangle$ state results are reclassified as $|1\rangle$ in post-processing. ESP can protect qubits against T_1 processes during readout, which can include readout-induced decay. We utilize ESP to improve qubit readout, and calibrate readout amplitudes to maximize readout fidelity with ESP turned on. In Tables A2 — A5, we plot the individual and simultaneous single-qubit confusion matrices with and without ESP. We note that some $|1\rangle$ state readout fidelities are very low without ESP. This is due to the fact that maximizing the readout fidelity with ESP sometimes results in higher readout amplitudes for better state separation; these higher amplitudes can result in readout-induced decay from $|1\rangle \to |0\rangle$. However, ESP provides intrinsic protection from these effects, since direct $|2\rangle \to |0\rangle$ decay is suppressed, thus providing a net benefit.

Even with improved readout fidelities using ESP, qubits can experience readout crosstalk during measurement. In Fig. A1b, we plot the frequency spectrum of the readout resonators for all eight qubits. We observe that several readout resonators are close in frequency. For example, the readout resonators for Q0 and Q1 are within ~4 MHz of each other, and the readout resonators for Q3, Q4, and Q5 are all within ~11 MHz of each other. Readout crosstalk can lead to context-dependent readout errors, in which the error on one qubit depends on the state of another qubit. This effect is apparent in the results presented in Fig. 1, in which the $|010\rangle$ and $|110\rangle$ states had drastically worse readout fidelities than the other preparation states.

(a) Qubit frequency spectrum

(b) Readout resonator spectrum

(c) Qutrit state discrimination
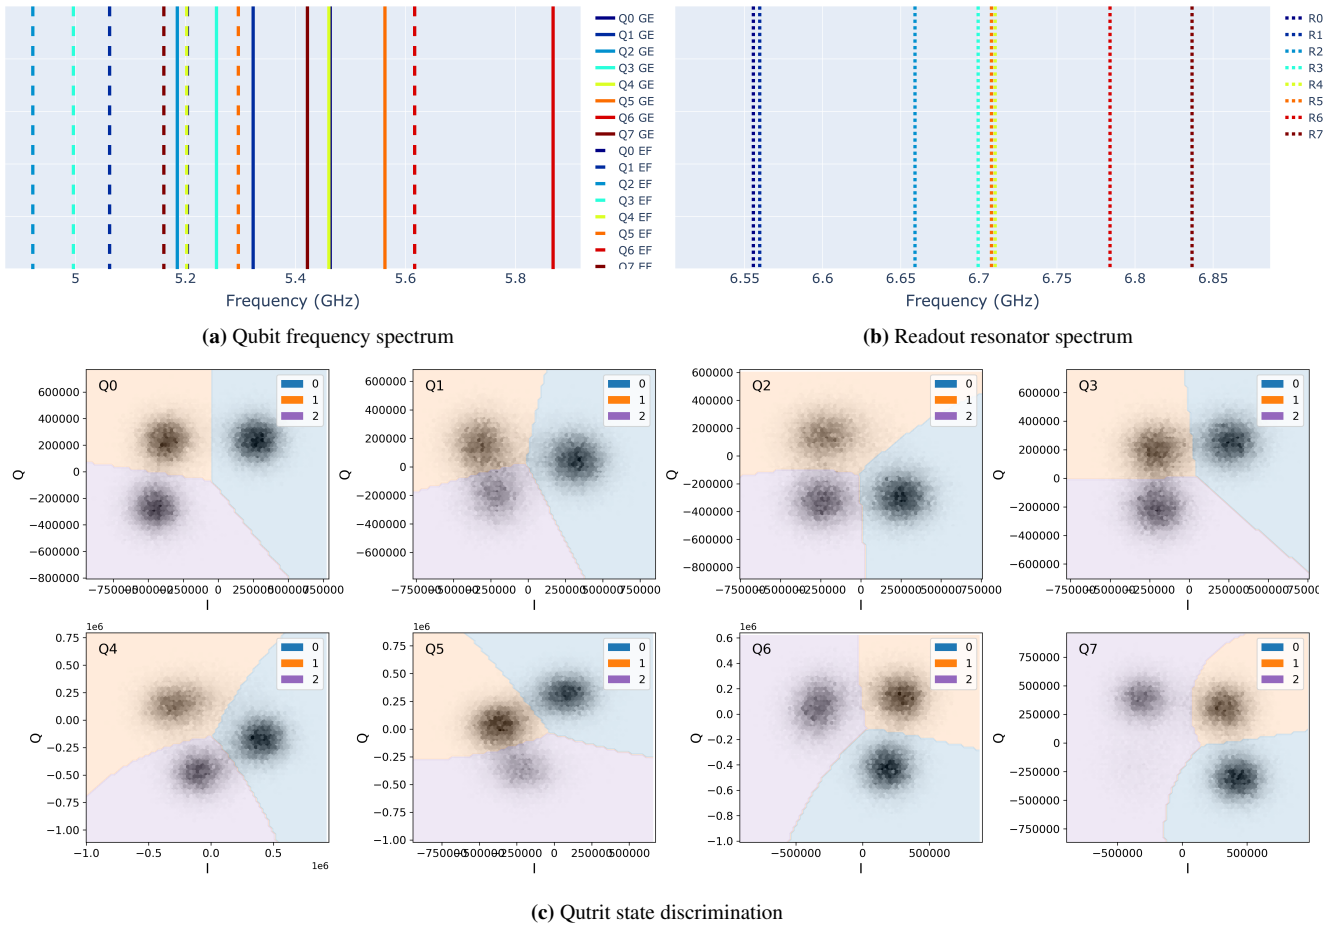
Figure A1: Qubit & Readout Characterization. (a) Frequency spectrum of the GE (solid lines) and EF (dashed lines) transitions of the 8 qubits on the quantum processor. (b) Frequency spectrum for the readout resonators coupling to the qubits. (c) Qutrit state discrimination is supported for all qubits on the quantum processor.

	Q0	Q1	Q2	Q3	Q4	Q5	Q6	Q7
T_1 (μs)	74.0(6.8)	129.0(3.5)	154.0(3.4)	49.7(1.1)	134.0(3.7)	82.0(2.5)	53.5(1.2)	19.9(0.2)
T_{2E} (μs)	36.0(1.9)	39.0(2.6)	77.0(3.2)	40.0(1.9)	50.0(2.2)	47.0(2.2)	40.0(1.9)	28.0(2.1)
RB iso. (10^{-3})	1.1(1)	0.75(4)	1.7(3)	0.87(8)	0.65(3)	0.56(3)	2.5(3)	4.6(5)
RB sim. (10^{-3})	5.3(9)	2.7(2)	2.9(3)	3.1(2)	7.0(5)	5.7(3)	3.4(3)	7.6(9)
CB sim. (10^{-2})				4.5(2)				

Table A1: Qubit characterization. Qubit coherence times (T_1 and T_2E) are listed above. The process infidelities for isolated and simultaneous single-qubit randomized benchmarking (RB) are listed as well. The average process infidelity of a cycle of simultaneous Paulis is measured using cycle benchmarking (CB).

	Q0	Q1	Q2	Q3	Q4	Q5	Q6	Q7
$P(0|0)$	0.989(2)	0.981(2)	0.968(3)	0.975(2)	0.967(3)	0.988(3)	0.978(3)	0.985(3)
$P(1|1)$	0.509(8)	0.918(4)	0.941(4)	0.476(6)	0.893(5)	0.617(6)	0.936(5)	0.909(5)

Table A2: Individual readout fidelities for all qubits without excited state promotion. $P(0|0)$ is the probability of measuring 0 when the qubit is prepared in $|0\rangle$, and $P(1|1)$ is the probability of measuring 1 when the qubit is prepared in $|1\rangle$. $P(1|1)$ is abnormally low for some qubits because the readout amplitudes have been optimized for use with ESP.

	Q0	Q1	Q2	Q3	Q4	Q5	Q6	Q7
$P(0|0)$	0.989(2)	0.982(3)	0.966(3)	0.976(3)	0.966(2)	0.987(2)	0.980(3)	0.986(2)
$P(1|1)$	0.967(3)	0.949(4)	0.969(4)	0.982(2)	0.912(6)	0.872(3)	0.982(3)	0.989(3)

Table A3: Individual readout fidelities for all qubits with excited state promotion. $P(1|1)$ has improved for all qubits compared to when ESP is not used. For some qubits, $P(1|1) > P(0|0)$.

	Q0	Q1	Q2	Q3	Q4	Q5	Q6	Q7
$P(0|0)$	0.998(2)	0.984(4)	0.968(3)	0.971(3)	0.973(2)	0.968(3)	0.974(4)	0.983(3)
$P(1|1)$	0.883(7)	0.838(7)	0.82(4)	0.879(3)	0.913(6)	0.855(9)	0.901(6)	0.886(5)

Table A4: Simultaneous readout fidelities for all qubits without excited state promotion. We observe that readout crosstalk can improve the readout fidelity of some qubits compared to the values in Table A2.

	Q0	Q1	Q2	Q3	Q4	Q5	Q6	Q7
$P(0|0)$	0.989(2)	0.976(2)	0.965(4)	0.964(5)	0.967(4)	0.956(3)	0.975(3)	0.985(2)
$P(1|1)$	0.973(8)	0.940(8)	0.968(3)	0.975(1)	0.947(5)	0.927(6)	0.983(3)	0.973(2)

Table A5: Simultaneous readout fidelities for all qubits with excited state promotion. $P(1|1)$ has improved for all qubits compared to when ESP is not used. For some qubits, $P(1|1) > P(0|0)$.